

Maximizing Expected Profit (good) vs. Maximizing Profit in the Expected World (bad)

Well-diversified shareholders will want the managers at the companies in which they've invested to make decisions which maximize expected return on investment. Formally, they will want managers to solve this problem:

$$\max_{\text{all policies}} E[\text{Return_On_Investment}(\text{policy}, \text{randomness})] .$$

Less formally, they will want managers to adopt policies which do well, on average, across all of the scenarios Nature might throw at them.

Example: In the World Airlines case, since bumping a passenger is less costly than flying with an empty seat, the optimal policy is to sell 365 tickets for the flight. This will lead to relatively more instances of bumping (the cheaper alternative) and fewer instances of empty seats (the more expensive alternative) – a desirable trade-off.

Unfortunately, all too often in today's business world decisions are actually made by focusing only on what would work best in the “expected” world. Formally, policies are chosen to solve this problem:

$$\max_{\text{all policies}} \text{Return_On_Investment}(\text{policy}, E[\text{randomness}]) .$$

Example: In the World Airlines case, the expected number of no-shows is 50, and the optimal policy in a world with (certainly) 50 no-shows is to sell 350 tickets. Not only is this policy not truly optimal across all possibilities, but it *completely* ignores the relative costs of bumping and flying with empty seats.

The lesson is simple: Don't wash the uncertainty out of a decision problem too soon!